On the Threshold Purchasing Behavior of Customers Facing Dynamically Priced Perishable Products

Yong-Pin Zhou · Ming Fan · Minho Cho
Department of Management Science
University of Washington Business School,
Seattle, WA 98195-3200
yongpin@u.washington.edu · mfan@u.washington.edu · mc9@u.washington.edu

August 26, 2005

Abstract

Dynamic pricing is a standard practice that firms use for revenue management. With the vast availability of pricing data on the Internet, it is possible for consumers to become aware of the pricing strategies used by firms and to develop strategic responses. In this paper, we study the strategic response of customers to dynamic prices for perishable products. As price fluctuates with the changes in inventory and the elapse of time, a strategic customer may choose to postpone a purchase in anticipation of lower prices in the future. We analyze a threshold purchasing policy for the strategic customer, and conduct numerical studies to study its impact on both the strategic customer's utility and the firm's revenue. We find that in most cases the policy can benefit both the strategic customer and the firm. In practice, the firm could encourage customers waiting by adopting a target price purchasing system.

Key words: revenue management; strategic customer response; dynamic pricing; threshold policy

1. Introduction

Pricing has been an age-old management issue, especially for perishable products facing uncertain demand. Under the common fixed-price scheme, if the price is set too low, potential revenue will be lost; and if the price is set too high, demand will be low and perishable products may be wasted when they expire. Revenue management (*a.k.a.* yield management) has become an increasingly popular management tool in selling perishable products. It is widely used not only in the hospitality industry (e.g. airlines, hotels, and cruise lines), but also in many other industries where products or capacity are perishable (e.g., golf course reservation, natural gas pipeline reservation, concert and ball game ticket sales, and fashion products). See McGill and van Ryzin (1999) for details. Essentially, revenue management is a method that aims to sell the right inventory unit to the right customer, at the right time and for the right price (Kimes 1989). This is mainly achieved through dynamic pricing and inventory allocation. More recently, firms begin to take advantage of the Internet to sell perishable products online (Choi and Kimes 2002, Liddle 2003). To many firms, the Internet offers a new opportunity to implement revenue management techniques such as dynamic pricing because price changes are easy, inexpensive, and potentially more effective.

Most of the research on revenue management focuses on developing optimal pricing and inventory allocation policies (e.g. Littlewood 1972, Belobaba 1989, Gallego and van Ryzin 1994, Zhao and Zheng 2000). These models generally assume that consumers are myopic and will purchase the products when the prices are below their reservation values. In contrast, our research examines how consumers strategically respond to the firm's dynamic prices over time. The growing use of the Internet provides an opportunity for consumers to gather information on companies' pricing policies and respond strategically. Since price variation for consumer products on the Internet is high, comparison-shopping can provide real benefits. The primary online shopping tools that consumers use today are shopbots that do price comparisons (Montgomery et al. 2004). Traditional shopbots compare prices spatially by checking the prices at various websites at roughly the same time. They do not anticipate possible future price changes at each site. More recently, researchers have shown interests in developing shopbots that can also compare prices temporally. One example is the "Hamlet" program that studies past trends in the variation of airline fares and establishes the patterns, which can then be used to decide whether the customer should make a purchase immediately or wait for possible future price reductions (Etzioni et al. 2003). Many critics doubt the effectiveness and the accuracy of Hamlet's prediction, however, because the underlying factors that determine the prices are not clearly understood (Knapp 2003).

In this research, we examine the behavior of the customer who strategically responds to dynamic prices, by timing the purchase. Since the price of a perishable product changes continuously over time, there is a chance to get the same product at a lower price by waiting. On the other hand, there is a chance that the product, which is available and affordable now, may be sold out or become unaffordable later if the strategic customer waits. We develop a threshold purchasing policy that balances this tradeoff, and examine its impact on both the customer and the firm. Focusing on the main factors that influence price, our analytical approach gives simple and effective solutions, and allows us to derive insights. We find that the strategic customer delay could benefit both the customer and the firm.

The rest of the paper is organized as follows. In the next section, we provide a brief review of the related literature. In Section 3, we develop the model, derive a threshold policy that helps the strategic customer to decide when to purchase, and use simulations to evaluate its benefit to the strategic customer. In Section 4, we numerically examine the impact of the threshold policy on the firm's revenue. We make the concluding remarks in Section 5.

2. Literature Review

There is an extensive body of literature on revenue management, mostly in the context of airline ticket sales. For a comprehensive review, see Talluri and van Ryzin (2004). Two different approaches complement each other. The first assumes that customers can be categorized into different classes (e.g. leisure and business travelers) and focuses the analysis on the allocation of capacity among these classes. Based on the demand forecast for each customer class, a "booking limit" of perishable products (e.g. airplane tickets) is computed for each customer type. These thresholds can vary over time as demand unfolds (Littlewood 1972, Brumelle and McGill 1993, Robinson 1995). The second approach focuses more on the dynamic pricing aspect of revenue management. Gallego and van Ryzin (1994) analyze the dynamic pricing policy for one type of product and homogeneous customers. The customers arrive randomly and their valuations for the product are also random. Important monotonicity properties are derived for the firm's optimal pricing policy. Zhao and Zheng (2000) extend this model to include non-homogeneous demand. Customers are time sensitive so their reservation price distribution may change over time. None of these models consider customers' reaction to the firm's pricing strategy, however. In their review of dynamic pricing models, Bitran and Caldentey (2003) point out "incorporating rationality on the behavior of customers" as an interesting field of research.

Rather than assuming customers are price takers, some marketing researchers have studied rational shopper behavior in the face of random price variations. In Ho, Tang and Bell's rational shopper model (1998), the firm chooses one of a finite set of pricing scenarios and rational shoppers react by purchasing more when the price is low and less when the price is high.

They find, among other things, that when price variability is high, the rational shoppers shop more frequently and buy fewer units every time. The type of product under consideration is the daily consumer product, which needs to be purchased and consumed repeatedly and continually over time. Consequently the main tradeoff for a rational shopper is between the purchase costs and the inventory holding costs. This differs from the one-time purchase of perishable products, which is the focus of this paper. Moreover, the price variation in Ho, Tang and Bell (1998) is random, while the price variation in our paper follows certain optimally determined curves that are given in the revenue management literature.

Besanko and Winston (1990) study a game between a monopolist, who sets prices for a new product over time, and strategic consumers, who decide whether to purchase now for the sure utility or postpone the purchase so as to maximize the future expected utility. In the equilibrium, the monopolist systematically reduces price over time. Elmaghraby et al. (2005) also study a game where the seller changes price over time, and the buyers submit the desired quantities at any given price. Liu and van Ryzin (2005) study a similar problem in a discrete time-period setting, and allow customers to be risk averse. All these three papers are based on the assumption that all consumers are present at the beginning of the game, which results in certain monotonicity properties. In our paper, customers arrive randomly over time. Therefore, the optimal price trajectory depends on the realization of the customer arrival process, and it may experience gradual decrease over time and sudden increase right after a purchase is made.

Aviv and Pazgal (2003) also study the strategic customer reaction to price variations, and allow customers to arrive over time. When forward-looking customers have information about future price discounts, they may decide to postpone their purchases to a later time when discounts are offered. In their model, there are only a pre-fixed number of price changes, and the price-setting firm announces the prices and the price change times ahead of time. While this may represent a retail-type environment, it does not apply to the situations where the firm continuously changes its price in response to the realization of stochastic demands. In addition, in Aviv and Pazgal (2003), the customers' valuations are homogeneous and decrease over time according to a deterministic function known to the firm. Thus, a customer who arrives at a certain time will have a deterministic valuation for the product. In contrast, our model assumes heterogeneous customer valuations and random customer arrivals. Consequently, the prices customers face are also stochastic.

Anderson and Wilson (2003) study consumer reactions to the dynamic allocation of airline seats to various fare classes. When all the low-price fares are closed, consumers may decide to wait before purchasing a ticket in the hope that a low-price fare class will reopen. The paper does not model consumer behavior explicitly, however.

The model in our paper is most closely related to the dynamic pricing model in Gallego and van Ryzin (1994). Gallego and van Ryzin (1994) assume that all customers are price-takers: those who can afford the price purchase right away and those who can't, leave. In our model, there is a strategic customer who is patient and would not purchase until the desired time or price of the product is reached. We are most interested in the effect of such a strategy on the customer's utility and the firm's total revenue.

The threshold purchasing policy derived from our model is related to the limit order trading in finance. Traders in financial markets can choose to place market orders or limit orders. While market orders are executed immediately at the market price, limit orders will wait to be executed at a predetermined price. To a certain extent, the regular customer behavior studied in traditional revenue management literature resembles the market order trading, while the strategic customer behavior studied in this paper resembles the limit order trading. There is an extensive literature in finance on the limit order trading. Harris and Hasbrouck (1996) find that limit orders placed at or better than the prevailing quote perform better than market orders. Chung et al (1999) find that more investors use limit orders when the bid-ask spread is wide. Forcault (1999) models the price formation and order placement decisions in the dynamic trading market. His main finding is that when price volatility increases, the market order trading becomes more costly and more traders find it optimal to submit limit orders. In terms of the impacts on the overall market, it is well recognized that limit orders play a significant role in supplying market liquidity and narrowing bid-ask spreads (Chung et al. 1999). The insights of the market microstructure literature prove valuable in helping us to interpret our results.

3. Strategic Customer Behavior

3.1 Dynamic Pricing Model

We assume that the firm's pricing strategy follows that in Gallego and van Ryzin (1994), which we call the GVR model. Therefore, we begin with a brief review of the GVR model. There is a fixed number, n, of one type of perishable product to be sold during a finite time horizon T. The product is perishable so all units left at the end of the sales period are worth nothing[1]. Let k denote the number of products left, $0 \le k \le n$, and t the time units left in the sale horizon, $0 \le t \le T$. As is the convention, t gets smaller as time goes by. Therefore, the state of the system can be described by the vector (k,t).

Customer purchases follow a price-sensitive Poisson process. That is, if price is p, then the instantaneous Poisson arrival rate is $\lambda(p)$, where $\lambda(p)$ is decreasing in p and $\lim_{p\to\infty} \lambda(p) = 0$. There is another way to interpret this customer arrival process: Let the arrival of all potential

[1] It will be straightforward to include an end-of-horizon salvage value for each unsold unit.

customers follow a Poisson process with a *constant* rate $\lambda(0)$. Moreover, let each customer's valuation of the product, *v*, have the cumulative probability distribution (CDF) $F(p) = 1 - \frac{\lambda(p)}{\lambda(0)}$. Thus, when the price is p, an arriving customer can afford the product with the probability of $1 - F(p)$. Consequently, the price-sensitive purchase arrival process is Poisson with the instantaneous rate $\lambda(0)\left[1 - F(p)\right] = \lambda(p)$. In this paper, the second interpretation will be used.

In any state (k,t), the firm chooses the best price p - or equivalently $\lambda(p)$ - to maximize its total expected revenue $J(k,t)$. Gallego and van Ryzin (1994) show that $J(k,t)$ is determined by the following equation, with boundary conditions $J(n,0)=0$ and $J(0,t)=0$:

$$\frac{\partial J(k,t)}{\partial t} = \sup_{\lambda}\left[\lambda p(\lambda) - \lambda\big(J(k,t) - J(k-1,t)\big)\right], \qquad \forall n \geq 1,\ \forall t > 0,$$

where $p(\lambda)$ is the inverse function of $p(\lambda)$.

The firm's dynamic pricing strategy is thus summarized in a pricing function $p(k,t)$. Clearly, it is reasonable to expect the price to be higher when fewer units are left, or when more time is left. Thus, $p(k,t)$ is decreasing in k but increasing in t. Those properties are proved in Gallego and van Ryzin (1994).

In the GVR model, all customers are price takers, which we call regular customers. In this paper, we assume there are two types of customers: regular customers (RC) and a strategic customer (SC). We assume that the strategic customer, with the help of software agents, is able to collect information about the firm's pricing policy $p(k,t)$ and the demand arrival function $\lambda(p)$, and use them to optimize the timing of her purchase. The strategic customer exhibits two major differences in her behavior from that of a regular customer. First, when a regular customer cannot afford the item, she simply leaves. In contrast, a strategic customer chooses to wait so that, if the price drops later, she may afford it. Second, when a regular customer can afford the item at the current price, she purchases right away. In contrast, the strategic customer may decide to postpone the purchase so that she may purchase the product at a lower price later.

3.2 Threshold Purchasing Policy

When the strategic customer arrives to find the system in the state (k,t), the decision for her is whether to purchase at the current price $p(k,t)$ or to wait. If she decides to wait, then what is the desired time or the price level to make the purchase? We study the following two policies:

THRESHOLD TIME POLICY (TTP). *With k products available, purchase if and only if there are* t_k *time units or less left in the sales time horizon. That is, purchase if and only if* $t \le t_k$.

THRESHOLD PRICE POLICY (TPP). *With k products available, purchase if and only if the price is below a certain threshold price level* p_k.

Since the strategic customer knows that the firm changes price dynamically, waiting a little bit to purchase may result in a lower price. How long she waits will have to depend on both the number of units left, k, and the time left, t. This results in the TTP. From another perspective, the strategic customer waits till a target price is reached, which is the TPP. Below, we show that the two policies are equivalent.

PROPOSITION 1. *The threshold price policy (TPP) is equivalent to the threshold time policy (TTP).*

PROOF. Because the pricing curves $p(k,t)$ are strictly increasing in t for the fixed k, it is easy to show that there exists a one-to-one relationship between t_k and $p_k = p(k,t_k)$ such that $t > t_k \Leftrightarrow p > p_k$, and $t < t_k \Leftrightarrow p < p_k$. Thus, there is a one-to-one correspondence between the TPP and the TTP. ■

Because the TTP and the TPP are equivalent, in this paper we will use them interchangeably.

If the strategic customer arrives with little time but many products left (i.e., small t and big k), the price may already be lower than her target price p_k so the strategic customer will purchase right away. In other situations, the strategic customer may wait. Clearly, during the wait, it is possible that another customer may arrive and make a purchase. In this case, k becomes $k-1$, and the strategic customer will continue to follow the above policies and wait for t_{k-1} (or p_{k-1}).

Now we study how the strategic customer determines the t_k s and p_k s. Let her have a valuation of v for the product. The objective for her is to maximize her utility, which is defined to be the difference between v and the price paid for the product. Clearly, the strategic customer will only purchase the product if the price is no more than v (i.e., no negative utility). If the strategic customer ends up unable to purchase the product because the price is higher than v, we say that the customer receives a utility of 0.

At any time, if the price is below the strategic customer's valuation, she has two actions: purchase now and get the sure utility, or wait till later to either get the product at a lower price or see the price jump due to other customers' purchases. The strategic customer must carefully balance the consequences of the two actions. We let the threshold t_k be the point at which the SC is indifferent between purchasing now and waiting a little longer[2].

PROPOSITION 2. *Let* t_k *be the solution to*

[2] That the threshold t_ks are such indifference points is intuitive. Moreover, due to bounded rationality, it is reasonable to assume that the strategic customer only considers these two options. A more rigorous approach would also consider the option of waiting to purchase at a more distant future time. In this case, even though we believe Proposition 2 still holds, we can only prove it for some special cases. Even when equation (1) is considered to be a heuristic, our numerical results in Section 3.4 show it is very effective.

$$\min\left\{p(k-1,t),v\right\}=p(k,t)+\frac{\frac{\partial}{\partial t}p(k,t)}{\lambda\left(p(k,t)\right)}. \tag{1}$$

If $t_k \geq t$ *, the strategic customer will purchase right away; and if* $t_k < t$*, the strategic customer will wait and the target purchase time is* t_k *.*

PROOF. Suppose that the strategic customer arrives in the state (k,t) and sees the price $p(k,t)$. We denote $q_i(k,t,\Delta t)$ the probability of i customers arriving during $[t,\ t-\Delta t]$ who can afford $p(k,t)$. It is easy to see $\lim_{\Delta t\to 0} q_0(k,t,\Delta t)=1$, $\lim_{\Delta t\to 0}\frac{q_1(k,t,\Delta t)}{\Delta t}=\lambda\left(p(k,t)\right)$, and $\sum_{i\geq 2} q_i(k,t,\Delta t)=o(\Delta t)$. If the strategic customer purchases the product at t, the realized utility is $v-p(k,t)$. If the strategic customer waits and purchases after Δt, the expected utility is

$$q_1(k,t,\Delta t)\max\left\{0,v-p(k-1,t-\Delta t)\right\}+q_0(k,t,\Delta t)\left[v-p(k,t-\Delta t)\right]+o(\Delta t)\,. \tag{2}$$

At the threshold t_k, the SC is indifferent between purchasing and waiting a little bit. By equating these two utilities and letting Δt go to 0, we obtain

$$\begin{aligned}&\lim_{\Delta t\to 0}\left[\frac{1-q_1(k,t,\Delta t)-q_0(k,t,\Delta t)}{\Delta t}\right]\left[v-p(k,t)\right]\\ &=\lim_{\Delta t\to 0}\left\{\frac{q_1(k,t,\Delta t)}{\Delta t}\left[p(k,t)-\min\left\{p(k-1,t-\Delta t),v\right\}\right]+q_0(k,t,\Delta t)\frac{\left[p(k,t)-p(k,t-\Delta t)\right]}{\Delta t}+\frac{o(\Delta t)}{\Delta t}\right\}.\end{aligned}$$

This amounts to $0=\lambda\left(p(k,t)\right)\left[p(k,t)-\min\left\{p(k-1,t),v\right\}\right]+\frac{\partial}{\partial t}p(k,t)$. Therefore, the time threshold t_k satisfies: $\min\left\{p(k-1,t),v\right\}=p(k,t)+\frac{\frac{\partial}{\partial t}p(k,t)}{\lambda\left(p(k,t)\right)}$.

■3.3 Exponential Valuation of the Customers

Equation (1) can be used to derive the thresholds for any given price strategy $p(k,t)$. To evaluate its efficiency, we will apply it to the case in which v follows an exponential distribution. This is the same distribution used in Kincaid and Darling (1963) and Gallego and van Ryzin (1994).

Let the arrival of potential customers follow a Poisson process with a constant rate of a. Each customer's valuation of the product, v, follows an exponential distribution with a rate of α. Consequently, when the price is $p(k,t)$, the probability that an arriving customer has a valuation v higher than $p(k,t)$ is $e^{-\alpha p(k,t)}$. Hence, the price-sensitive Poisson arrival rate is $\lambda(p(k,t)) = ae^{-\alpha p(k,t)}$. For simplicity of notation, α is set to one.

Under these assumptions, Gallego and van Ryzin (1994) show that the optimal pricing policy for the firm satisfies:

$$p(k,t) = J(k,t) - J(k-1,t) + 1 \tag{3}$$

where $J(k,t)$ is the maximum revenue function for the firm and it satisfies:

$$\frac{\partial J(k,t)}{\partial t} = \lambda(k,t) \tag{4}$$

and

$$J(k,t) = \log\left(\sum_{i=0}^{k} (\frac{at}{e})^i \frac{1}{i!}\right). \tag{5}$$

In what follows, we will further characterize these functions and derive properties that will simplify the analysis of Equation (1). To streamline the exposition, we will use $\lambda\left(k,t\right)$ instead of $\lambda\left(p(k,t)\right)$, and define $g(k,t) = p(k,t) + \frac{\frac{\partial p(k,t)}{\partial t}}{\lambda\left(k,t\right)}$. The proofs of the following two lemmas can be found in the Appendix.

LEMMA 1. $p(k,t)+\dfrac{\frac{\partial p(k,t)}{\partial t}}{\lambda(k,t)}$ *is increasing in t.*

LEMMA 2. $p(k-1,t)>p(k,t)+\dfrac{\frac{\partial p(k,t)}{\partial t}}{\lambda(k,t)}$.

Based on Lemmas 1 and 2, we can find the optimal purchasing thresholds.

PROPOSITION 3. *(i) The t_k s for the TTP are solutions to* $v=p(k,t)+\dfrac{\frac{\partial p(k,t)}{\partial t}}{\lambda(k,t)}$.

(ii) A unique finite solution, t_k, exists for every k if and only if $v\geq 1$.

PROOF. Lemmas 1 and 2 imply that the LHS and the RHS of Equation (1) look like the graph displayed in Figure 1.

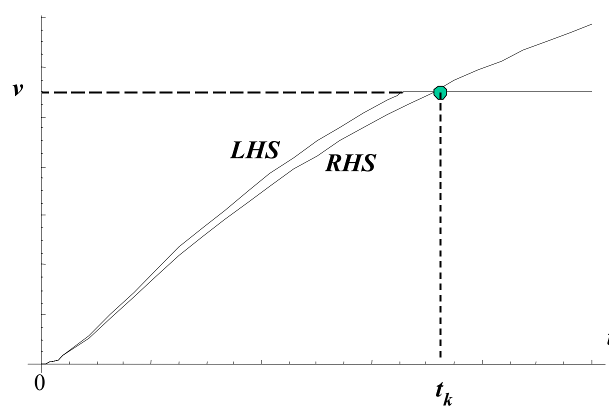

Figure 1. Solution to the Threshold Policy.

(i) From Figure 1, it is clear that, because $p(k-1,\ t)$ is always greater than the RHS, t_k is the intersection of v and the RHS. In effect, Equation (1) can be simplified to

$$v=p(k,t)+\frac{\frac{\partial}{\partial t}p(k,t)}{\lambda\big(p(k,t)\big)}. \tag{6}$$

(*ii*) Next, we prove that the two curves will intersect if and only if $v \geq 1$. Since the RHS is increasing in *t*, its minimum is achieved at $t = 0$, which is 1. So clearly *v* needs to be at least one. On the other hand, when *t* goes to 0, the RHS goes to 1; and when *t* goes to ∞, the value goes to ∞. Because the RHS is continuous, we conclude that for any $v \geq 1$, there exists a t_k such that the equality holds. The uniqueness follows easily from the monotonicity of the RHS (Lemma 1). ■

PROPOSITION 4. *The solution of the TTP has the following properties:*

(i) *The t_ks are increasing in v.*

(ii) *The t_ks are decreasing in a.*

PROOF. (*i*) follows immediately from Figure 1. For (*ii*), we note that

$$\begin{aligned}\frac{\partial g(k,t)}{\partial a} &= \frac{\partial p(k,t)}{\partial a} + e^{-[p(k-1,t)-p(k,t)]}\left[\frac{\partial p(k-1,t)}{\partial a} - \frac{\partial p(k,t)}{\partial a}\right] \\ &= \frac{\partial p(k,t)}{\partial a}\left[1 - e^{-[p(k-1,t)-p(k,t)]}\right] + e^{-[p(k-1,t)-p(k,t)]}\frac{\partial p(k-1,t)}{\partial a}.\end{aligned}$$

Because $p(k,t)$ is increasing in *a*, it follows that $g(k,t)$ is also increasing in *a*. Clearly from Figure 1, as the RHS increases and *v* stays the same, the intersection point, t_k, decreases. ■

Proposition 3 reduces the computational effort for the time thresholds and facilitates further theoretical analysis. It also has a simple interpretation: When the strategic customer decides not to purchase right now, two things are possible: the price may go up if another customer arrives and the effect of this is $\lambda(k,t)\left[p(k-1,t) - p(k,t)\right]$; or if there is no other arrivals then the price will gradually go down over time and the effect of this is $\frac{\partial p(k,t)}{\partial t}$. Proposition 3 shows that the first, price-jump effect always exceeds the second, time effect.

Therefore, if *v* is very high, the strategic customer will always purchase immediately. However, the existence of a finite *v* limits the first effect, and makes the waiting option more attractive. One can also easily deduce that the lower the v, the more restrictions it puts on the price-jump effect, and the customer is more willing to wait.

This is formally stated in Proposition 4. Intuitively speaking, when v is small, the utility for the strategic customer, if she purchases the product right away, is small; so the strategic customer has little to lose if she waits and other regular customers make purchases (and the price goes above v), but she has much to gain if the price keeps dropping. On the other hand, the strategic customer with a higher v will have a higher loss if the price jumps, but the same amount to gain by waiting as that with a lower v. As a result, the strategic customer with a higher v will purchase earlier. Numerically, this holds true especially for $k = 1$. For $k > 1$, the t_ks are quite insensitive to v.

Proposition 4 also states that the bigger the a, the smaller the t_k. This means that if the product is "hot", then the customer will want to wait longer. This seems counter-intuitive, but it makes sense because if the demand rate is high, the firm also knows it. As a result, for the same k and t, the price will be higher for a higher a. Therefore the utility to gain for a customer with a fixed v is lower. Therefore, by waiting, the strategic customer risk losing this current utility, but the loss is smaller now (since a is larger). So the customer is willing to wait longer.

When the strategic customer follows the TTP, she needs to estimate the following three parameters to determine her time thresholds:

- time left, t.
- number of products left, k
- arrival rate, a

Of course, t is usually easy to estimate. Thus, it remains to estimate k and a. In the airline industry, for example, some of the online booking sites provide information about how many tickets are still available (by showing the airplane layout and marking the available seats) *before* the customer purchases the ticket[3]. Therefore it is possible to get a good estimate of k as well. The most difficult parameter to estimate is a.

The following proposition shows that when a customer follows the TPP, there is no need at all to estimate the arrival rate a:

PROPOSITION 5. *The solution to the TPP,* p_k*, is independent of a.*

PROOF. First of all, note the following:

$$\begin{aligned} p_k = p(k,t_k) &= v - \frac{\frac{\partial p(k,t_k)}{\partial t}}{\lambda(k,t_k)} = v - \frac{\lambda(k,t_k) - \lambda(k-1,t_k)}{\lambda(k,t_k)} \\ &= v - 1 + \frac{\lambda(k-1,t_k)}{\lambda(k,t_k)} = v - 1 + e^{-[p(k-1,t_k) - p(k,t_k)]}. \end{aligned} \tag{7}$$

Because $J(k,t) = \log\left(\sum_{i=0}^{k} \frac{(at/e)^i}{i!}\right)$ and $p(k,t) = J(k,t) - J(k-1,t) + 1$, both $J(k,t)$ and $p(k,t)$ depend on a only through the product at. So if we let $x = at$ then both $J(k,t)$ and $p(k,t)$ become functions of only x (i.e. they are free of a). Therefore, we can simply solve (6) to obtain $x_k = at_k$ and plug them into (7) to compute p_k. The p_ks thus computed are all free of a. ■

From the proof we see that the p_ks are independent of a and the t_ks depend on a only through the product at_k. This makes sense because what's important for the strategic customer at time t, for a fixed inventory level k, is not the arrival *rate* of other customers, but rather the expected *number* of other customers who will arrive later. This is the product of the arrival rate and how much time is left, at.

[3] For example, visit Northwester, Delta, or American Airline's websites. Also, see "Cranky Consumer: Testing Out Airline Web Sites" by Sam Schechner in Wall Street Journal, Feb 15, 2005.

That the threshold prices can be determined with only *k* and *t* makes the TPP a lot easier to use. Also, it is worth noting that the use of the TPP is quite similar to that of the limit order in stock trading: a customer arrives to find the current prevailing market price and decides to transact later when a threshold price is reached. We will have more discussion on this later.

Table 1. Numerical Results, $n = 25$

	$a = 60$			$a = 70$		
	$v = 1.01$	$v = 1.03$	$v = 1.05$	$v = 1.01$	$v = 1.03$	$v = 1.05$
Time SC arrived	0.659	0.474	0.537	0.109	0.276	0.306
Time product bought	0.829	0.885	0.720	0.564	0.389	0.518
SC arrival price	1.040	1.152	1.065	1.108	1.045	1.112
SC purchase price	1.004	1.011	1.027	1.006	1.020	1.033
SC Utility	0.006	0.019	0.023	0.004	0.010	0.017

The example in Table 1 and Figure 2 illustrates how the threshold policy works for the strategic customer. Numerical results show that the threshold policy uses a different threshold time (t_k) for each inventory level (k). As shown in Figure 2, the threshold time decreases in k (this decrease is also observed in all the numerical tests we carried out for the

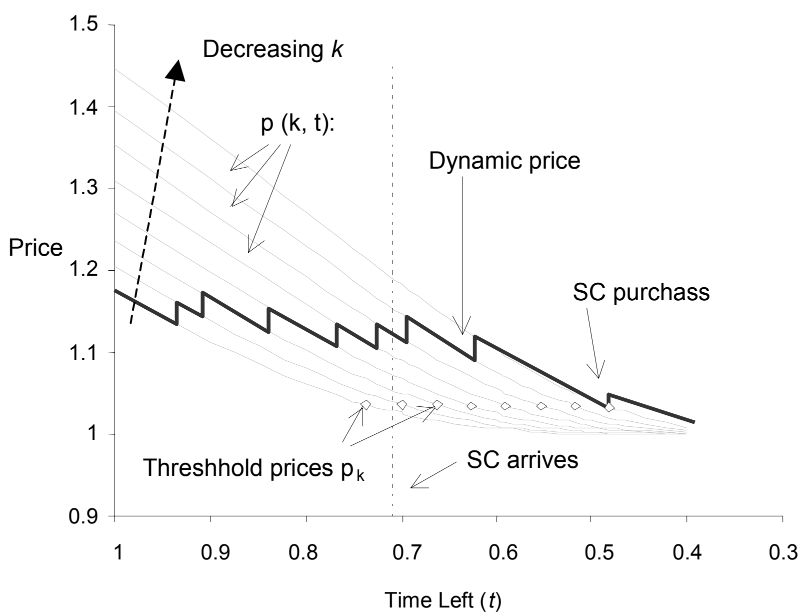

Figure 2. Strategic Consumer Purchasing Policy ($a = 70$, $n = 25$, $v = 1.05$)

simulations in the next section), which suggests that when inventory is higher, the strategic customer should wait shorter. The reason is that, when inventory is high, the firm's price will be low, which means the strategic customer does not need to wait long for the price to drop below the threshold level.

3.4 Benefits to the Strategic Customer

We conduct simulation studies to examine the benefit of using the TPP to the strategic customer. For every sample path of all customer arrivals, we run two simulations simultaneously. In Simulation 1, we randomly pick a customer to be the strategic customer. This strategic customer will follow the TPP. Simulation 2 is identical to Simulation 1 except that we replace the strategic customer with a regular customer. If her valuation is less than the current price, the regular customer in Simulation 2 will leave the market while the corresponding strategic customer in Simulation 1 will wait. If her valuation is higher than or equal to the current price, the regular customer in Simulation 2 will purchase the item right away, while the strategic customer in Simulation 1 may still wait depending on the TPP.

We compute two measures of the benefit to the strategic customer to follow the TPP. The first is the difference in utility gain between the strategic customer in Simulation 1 and her corresponding regular customer in Simulation 2. The second is the difference in the probability of obtaining the product between the two customers. We report the results in Figures 3 and 4 respectively.

It is worth noting that the price formula in Equation (3) yields a minimum price of 1. Therefore, any customer with valuation of less than 1 will never be able to afford the product. In the simulation tests, we allow the customer valuations to follow the exponential distribution, but will focus our attention only on the customers with valuations no less than 1.

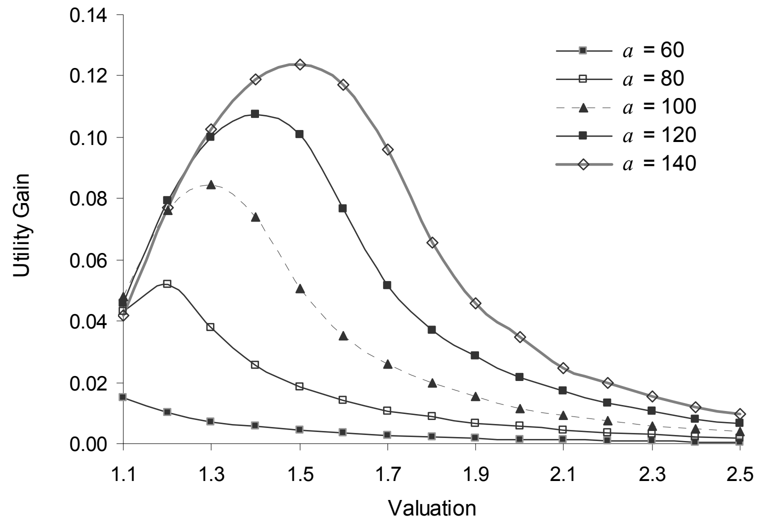

Figure 3. Average Utility Gain for SC

As shown in Figure 3, the strategic customer consistently outperforms the corresponding regular customer. The highest performance difference occurs when the strategic customer's valuation of the product is intermediate. Our explanation is that when the valuation is low, the maximum utility that can be obtained by the strategic customer is limited; so is the difference of utility between the two types of customers. When the strategic customer's valuation is high, her target prices will also be high. Oftentimes she will purchase the product immediately upon arrival. Thus, on average the strategic customer does not gain much utility than the corresponding regular customer. It is also interesting to note that the utility gain increases with the arrival rate. With a higher arrival rate, the average product price will also be higher. The strategic customer's benefit of using the TPP is higher under those situations.

For the strategic customer with a high valuation of the product, she will not improve her chance of getting the product by waiting because both the strategic customer and the corresponding regular customer are likely to afford the product. Therefore, it is expected that the improvement in the probability of getting the product mostly occurs when valuations are low. This is confirmed by results in Figure 4.

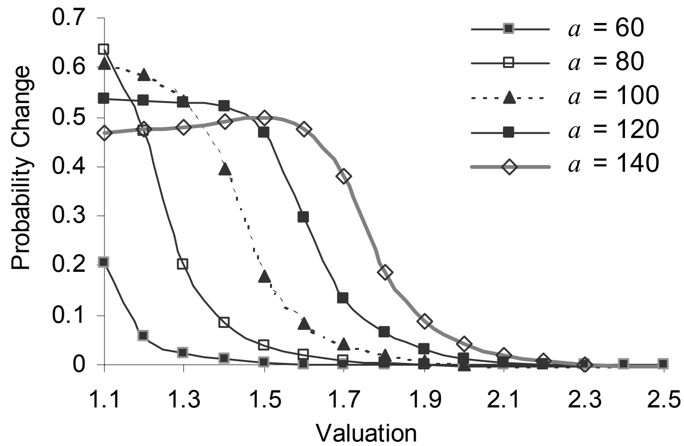

Figure 4. Percentage Point Increase in the Probability of Obtaining the Product for the SC

Figures 3 and 4 suggest that by following the TPP, the strategic customer may benefit because (1) she may improve her chance of getting the product if she could not afford it upon arrival, or (2) she may get a lower price later. The question is which effect is more dominant. To answer this, we perform more detailed analysis. We categorize all the simulation outcomes into two cases. In Case 1, the strategic customer cannot afford the product upon arrival. In Case 2, the strategic customer can afford the product upon arrival. The corresponding regular customer, who has the same valuation with the strategic customer, will get a utility of 0 in Case 1. Thus, the strategic customer will always be better off in Case 1. In Case 2, there are two possibilities. When the strategic customer manages to purchase the product, she will be able to obtain the item for a lower price and, thus, is better off by waiting. However, it is possible that the strategic customer waits but does not get the product while the corresponding regular customer purchases the product. Under this situation, the strategic customer is worse off.

Table 2 shows that the expected gain is always positive in both Cases 1 and 2, suggesting that on average the strategic customer is always better off. Furthermore, we find that the expected gain predominantly comes from Case 1 when the valuation is low and the arrival rate is high. In those situations, the price is less affordable and the TPP allows the strategic customer to have the chance to purchase the product at a price lower than his valuation. When the price is more affordable due to either a lower arrival rate or a higher consumer valuation, a higher percentage of the expected gain comes from Case 2.

Table 2. Percentage of the Expected Benefits from Case 1

	Arrival Rate				
Valuation	60	80	100	120	140
1.1	81.4	96.5	98.5	99.2	99.4
1.2	75.0	90.8	97.0	98.3	98.9
1.3	68.2	81.1	94.4	97.3	98.4
1.4	63.0	76.0	88.7	95.4	97.5
1.5	57.8	69.4	80.2	92.0	95.8
1.6	52.1	64.4	75.2	84.1	93.8
1.7	46.2	58.3	70.5	77.8	88.0
1.8	41.9	53.3	65.1	74.4	80.9
1.9	43.2	51.4	60.4	68.1	74.1
2.0	41.9	45.5	56.2	62.5	70.2
2.1	35.0	43.7	52.0	59.0	64.8
2.2	32.1	37.6	50.5	55.5	58.8
2.3	21.6	34.2	45.3	51.8	52.9
2.4	13.9	29.3	41.7	44.9	51.6
2.5	11.9	27.0	40.3	43.9	44.8

4. Firm Response

4.1 Impacts on the Firm

Since the use of the TPP benefits the strategic customer, one may expect that the firm will be worse off if it continues to use the original dynamic pricing policy. Revenue could decline because strategic customers will delay their purchases and pay lower prices. On the contrary, we find the firms by and large do better with strategic customers.

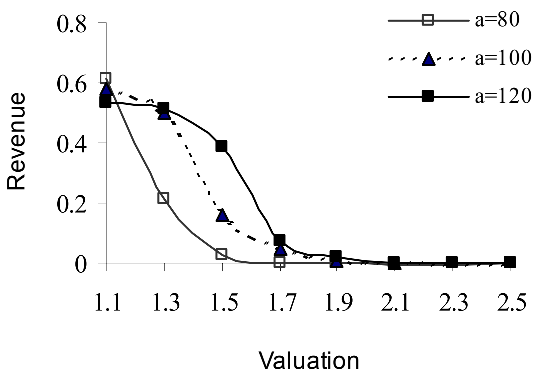

Figure 5. Revenue Increase

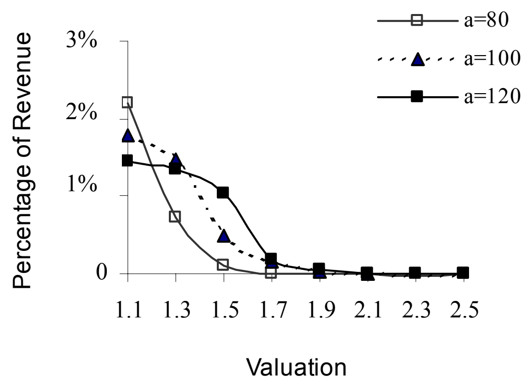

Figure 6. Percentage of Revenue Increase

From simulation results (Figures 5 and 6), we see that the impact on firm revenue is non-negative across all customer valuations. It is important to note that the increase in firm revenue is significant when the customer valuation is low (a 1% revenue increase can translate into a significant increase in profit). Since most likely the customer will have a low valuations because of her exponential valuation distribution, the firm's overall revenue increase is significant. We also see that the pattern of the increase in sales (number of tickets sold) is consistent with revenue increase, suggesting that sales increase is likely the major cause of revenue increase (Figures 7 and 8).

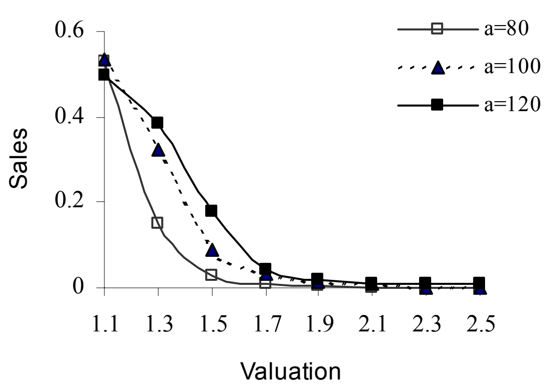

Figure 7. Sales Increase

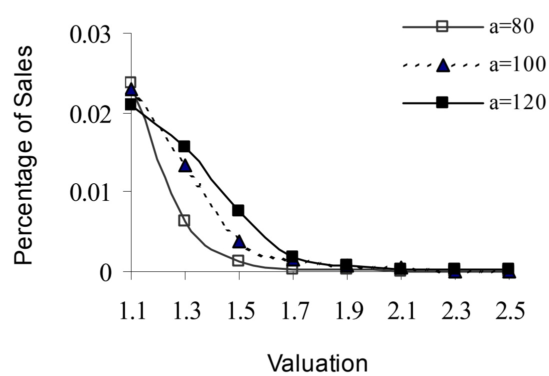

Figure 8. Percentage of Sales Increase

With a strategic customer delaying her purchase, there are two likely effects on firm revenue. First, the strategic customer may purchase at a lower price, or not purchase at all if the price increases beyond her valuation. This impact on firm revenue is negative. Second, when the strategic customer could not afford the item upon arrival, her waiting essentially keeps the demand, which would have otherwise been lost, in reserve. Consequently, the firm can both maintain a higher price and reduce the number of unsold items later on. This impact on firm revenue is positive. Our results suggest that the second, positive, effect dominates the first, negative effect.

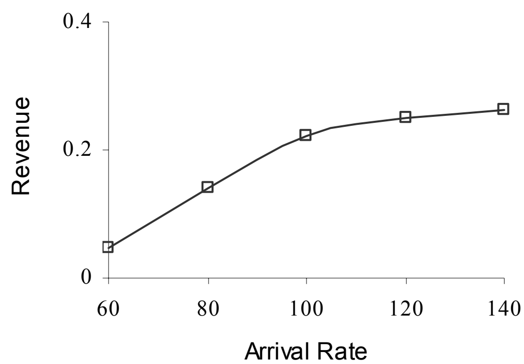

Figure 9. Average Revenue Increase by Arrival Rate

We conduct further simulations by following a random strategic customer, whose valuation follows the exponential distribution. We find the average firm revenue improvement increases in the arrival rate (Figure 9). The average sales increase also increases in the arrival rate with a similar pattern as Figure 9 and we omit the graph. Figures 3 and 9 together suggest that when the product is "hot" (higher demand relative to supply), the use of the TPP results in higher benefits for both the strategic customer and the firm. To explain this, we note first that when the arrival rate is high, the price is also high. Individual customers are more likely to be priced out of the market when they arrive. This may even happen to high-valuation customers at the beginning of the sales horizon. With the strategic customer's waiting, each time price drops below the threshold price of the strategic customer, the strategic customer will purchase. Thus, the strategic customer provides a valuable demand cushion for the firm, especially when the arrival rate is high.

Moreover, when the arrival rate is high, price volatility is also high. This is verified by numerical results in Figure 10, which examines the relationship between arrival rate and price volatility in the original GVR model. Using either the standard deviation or the coefficient of variation (CV) of the transaction prices in the GVR model as the measure of price volatility, we find price volatility indeed increases in the customer arrival rate. Intuitively, with a higher arrival rate, the firm will price products higher. However, if expected demand does not materialize, the firm has to reduce the price more sharply. Therefore, the higher arrival rate leads to higher price volatility.

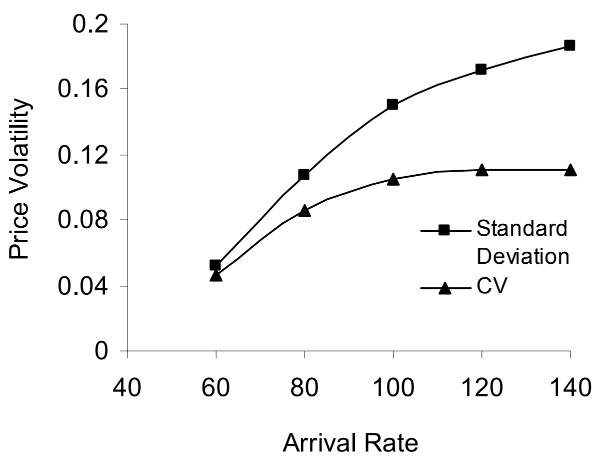

Figure 10. Price Volatility and Arrival Rate (n = 25)

It seems that the increase in both the strategic customer's utility and the firm's revenue can be explained by the increase in price volatility when the arrival rate is high. This is consistent with the results in financial literature on the limit order trading discussed earlier. In general, when the market is more volatile, individual market participants can benefit by being patient and waiting to a threshold price. Essentially, it is a form of transferring demand over time. When demand is stochastic, this practice will help improve the overall market performance as well. In financial market, limit orders narrow the bid-ask spread (Chung et al. 1999) and reduce transaction costs. In the market of perishable products studied in our research, the use of the TPP reduces wasteful inventory and increases firm revenue.

4.2 Option to Set Target Price

Realizing that the strategic customer's waiting could increase revenue, firms may develop a system to encourage customers' waiting. An interesting analogy, as mentioned early, is that allowing limit orders in financial markets helps to improve the overall market performance (Chung et al. 1999). There are many options to encourage customers to stay around rather than leave instantly when prices are too high. We consider a simple system that allows customers to indicate an intention of future purchase. For example, the firm can ask the customer to create a "wish list" of the product and the target price, as well as the email address where the customer can be informed when the price is reached. In this section, we numerically investigate the impact of such a system on the firm's revenue.

Such a target purchase system will be open to all the customers, but not everyone will use it. In our simulations, we let each customer choose to use such a system with a certain probability (we will continue to call them the strategic customers), and systematically vary this probability. Prior studies have reported that, even after many years in existence, online searching activities are still limited (Johnson et al. 2004, Montgomery et al. 2004). Therefore, we expect the use of such a target purchase system to be limited as well. Consequently, we vary the proportion of strategic customers from 2.5% to 15%.

Even among the strategic customers, it is likely that some will leave the system and purchase somewhere else or decide not to purchase at all before their target price is reached. Therefore, in our simulation we also vary the impatience ratio, which is the probability that a strategic customer will have left the system before her target price is reached. The impatience ratio reflects both the level of competition (e.g. how many airlines fly between the city pair on that date) and the level of customer loyalty (e.g. whether the customer belongs to a loyalty program). The higher the impatience ratio, the higher the probability the customer will leave. In our simulations we will also systematically vary the impatience ratio.

For simplicity, the firm allows a customer to leave only one target price[4]. Therefore, the strategic customers need to determine an inventory-independent target price. We use a simplified TPP heuristic for that. Using the simplified TPP heuristic, when a strategic customer arrives, she first estimates the average inventory the firm will carry from this time forward, and then use this average inventory and the TPP to calculate her single target price.

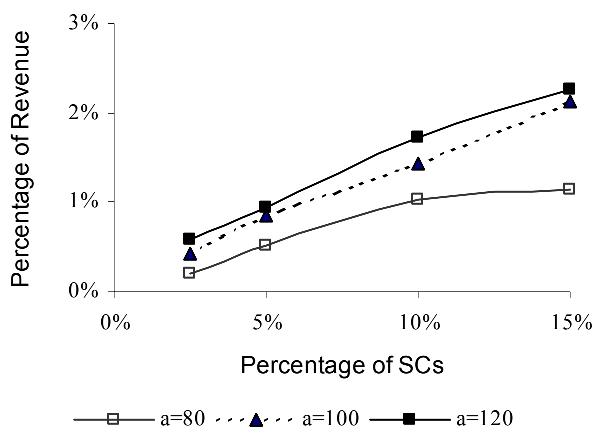

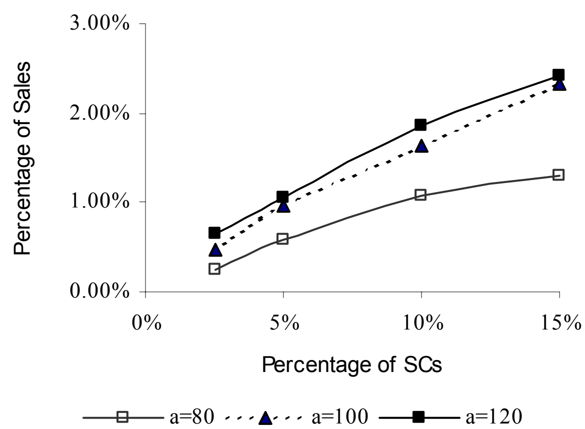

Figure 11. Revenue and Sales Increase under a Single Target Price. Impatience ratio = 10%.

---

[4] It is unlikely that the firm will allow the customers to leave target prices based on the inventory level. Theoretically, the customer can periodically check the inventory level and modify her target price accordingly. This calls for such a substantial amount of work on the customer's part that they will not use it frequently in practice. Therefore, to simplify analysis, we focus on the case where the customers leave one price and do not change it as time and inventory level change, except when leaving.

Figure 11 shows the impact of the strategic customer proportion on firm revenue and sales. It is clear that more strategic customers help the firm to increase its revenue more, as they provide a larger demand cushion for the firm so that the price does not drop too low[5]. This also explains the result that the firm will be able to sell more products when the strategic customer proportion increases.

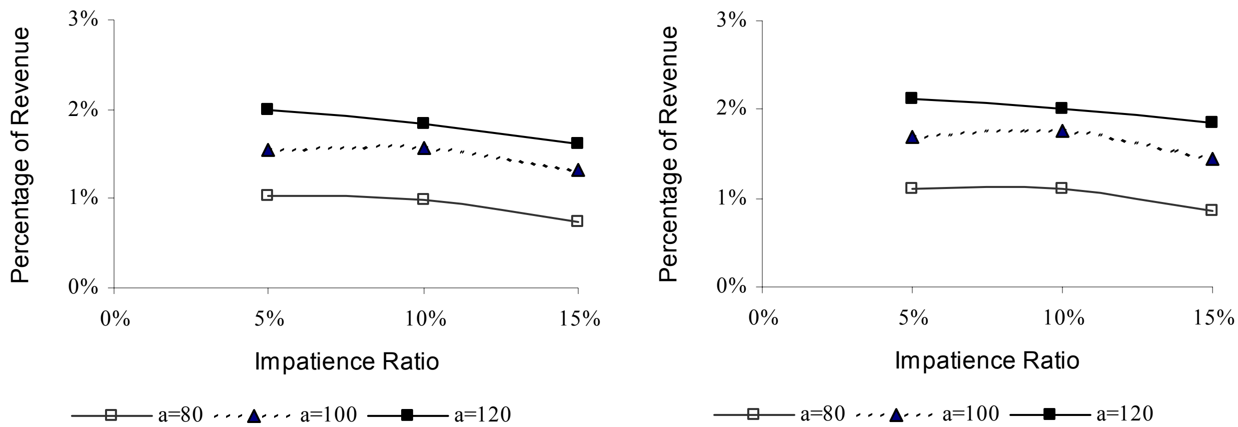

Figure 12. Percent of Revenue Change under a Single Target Price. SC Proportion = 10%

Figure 12 displays the impact of impatience ratio on firm revenue and sales. The results indicate that both will increase if customers are provided with such a purchase option. Not surprisingly, both the revenue and sales increase more when the impatience ratio is lower.

Figures 11 and 12 also reveal that the revenue increase is higher with a higher arrival rate. This again suggests that firms with more price volatility should have a stronger incentive to provide such a purchase option to their customers.

[5] Of course, the results are only valid in the range of strategic customer proportion. One can imagine that in the extreme when all customers are strategic (i.e. proportion is 100%), the system dynamic becomes completely different and a different analysis is needed. In practice, however, we believe our relatively low range of strategic customer proportion to be more likely.

5. Conclusions and Future Research

In this research, we study the strategic response of customers to dynamic prices of revenue management. The strategic customers wait to purchase at specific target prices (TPP) or target times (TTP) that depend on the customers' valuation of the product and the current inventory level. We conduct simulations to study the performance of the TTP/TPP. We show that customers benefit by following these policies. In particular we find that when the customer valuation is low or the arrival rate is high, most of the utility gains come from the improved probability of getting the product by waiting (hence, the utility improves to non-zero from zero). When the customer valuation of the product is high, then most of the benefit comes from the lower price by waiting. Overall, the benefit is the greatest for low-valuation customers (low v) and hot products (high a).

We also show that the firm also benefits from having strategic customers who follow the TTP/TPP. This result first seems to be counter-intuitive until one realizes that this is not a zero-sum game. The firm may benefit because, while regular customers will leave the market if they cannot afford the product upon arrival, strategic customers are kept in the waiting pool, especially early in the sales period when the product price is usually higher. This waiting of strategic customers provides a cushion to price volatility and prevents the price from falling too low. It also serves as an additional demand that helps to reduce wasted inventory at the end. This benefit is somewhat similar to the benefits of limit orders in stock trading which provide liquidity to the market (Chung et al. 1999, Forcault 1999).

When high-valuation strategic customers wait and get a lower price, this will negatively affect the firm's revenue. But our results show that, in general, the potential revenue loss from the delay of purchase is limited. High-valuation customers, as it turns out, have higher target prices, and are very likely to purchase immediately upon arrival. Keeping low-valuation customers in the waiting pool helps to reduce wasted inventory and prevent firms from deep price discounts toward the end. This could be especially beneficial to industries with a fixed cost for the products, e.g. airline tickets and hotel rooms, where the revenue loss of each wasted inventory is large.

This discovery of benefits to the firm is important as it encourages companies to develop systems that can allow customers to place a "limit order" for the products or services. Actual implementation can be flexible. Customers can choose to be notified of price changes through emails. A promising direction for future research is to model the impact of such practices on firm's revenue analytically and quantify the tradeoffs between the higher sales and the lower prices paid by some customers.

The analytical results in this paper are based on the modeling of a single strategic customer. Nevertheless, simulation results show that with multiple strategic customers, the benefits to each individual strategic customer (the utility gain and the probability of obtaining the product) remains quite stable, while the benefit to the firm (higher revenue) increases with the proportion of strategic customers. This result holds for small proportions of strategic customers (2.5%-15%). We conjecture that as the proportion increases, the benefit to the firm will plateau. Moreover, as strategic customers expect there to be other strategic customers waiting, their target times/prices will also adjust accordingly. How this interaction will affect the customers' purchase behavior, as well as the customer utility and firm revenue is a topic for future research.

Appendix

PROOF OF LEMMA 1.

From (3) and (4), $\frac{\partial p(k,t)}{\partial t} = \frac{\partial J(k,t)}{\partial t} - \frac{\partial J(k-1,t)}{\partial t} = \lambda(k,t) - \lambda(k-1,t)$. We also know that

$\lambda(k,t) = ae^{-p(k,t)}$. Therefore, because $g(k,t) = p(k,t) + \frac{\frac{\partial p(k,t)}{\partial t}}{\lambda(k,t)}$,

$g(k,t) = p(k,t) + 1 - \frac{\lambda(k-1,t)}{\lambda(k,t)} = p(k,t) + 1 - e^{-[p(k-1,t)-p(k,t)]}$.

Therefore,

$$\begin{aligned}\frac{\partial g(k,t)}{\partial t} &= \frac{\partial p(k,t)}{\partial t} + e^{-[p(k-1,t)-p(k,t)]}\left[\frac{\partial p(k-1,t)}{\partial t} - \frac{\partial p(k,t)}{\partial t}\right] \\ &= \frac{\partial p(k,t)}{\partial t}\left[1 - e^{-[p(k-1,t)-p(k,t)]}\right] + e^{-[p(k-1,t)-p(k,t)]}\frac{\partial p(k-1,t)}{\partial t}.\end{aligned}$$

Because $p(k,t)$ is increasing in *t* and decreasing in *k*, it follows that $\frac{\partial g(k,t)}{\partial t} \geq 0$. ■

PROOF OF LEMMA 2.

$\frac{\frac{\partial}{\partial t}p(k,t)}{\lambda(k,t)} = \frac{\lambda(k,t)-\lambda(k-1,t)}{\lambda(k,t)} = \frac{e^{-p(k,t)} - e^{-p(k-1,t)}}{e^{-p(k,t)}} = \frac{-e^{-\xi}\left[p(k,t)-p(k-1,t)\right]}{e^{-p(k,t)}} = \frac{e^{-\xi}\left[p(k-1,t)-p(k,t)\right]}{e^{-p(k,t)}}$ It is $< \left[p(k-1,t)-p(k,t)\right]$, for some $\zeta \in (p(k,t), p(k-1,t))$.

clear then $p(k-1,t) > p(k,t) + \frac{\frac{\partial}{\partial t}p(k,t)}{\lambda(k,t)}$. ■

Reference

Anderson, C.K. and J.G. Wilson. 2003. Wait or buy? The strategic consumer: Pricing and profit implications. *Journal of the Operational Research Society* 54 (3), 299-306.

Aviv, Yossi and Amit Pazgal. 2003. Optimal pricing of seasonal products in the presence of forward-looking consumers. Working Paper, Olin School of Business, Washington University.

Belobaba, Peter. 1989. Application of a probabilistic decision model to airline seat inventory control. *Operations Research* 37, 183-197.

Besanko, David and Wayne Winston. 1990. Optimal price skimming by a monopolist facing rational consumers. *Management Science* 36(5), 555-567.

Bitran, Gabriel and René Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing & Service Operations Management* 5 (3), 203 –229.

Brumelle, S. L. and J. I. McGill. 1993. Airline seat allocation with multiple nested classes. *Operations Research* 41 127-137.

Choi, Sunmee and Sheryl E. Kimes. 2002. Electronic distribution channel's effect on hotel revenue management. *Cornell Hotel and Restaurant Administration Quarterly* 43 (3), 23-31.

Chung, K., B. Van Ness, R. Van Ness. 1999. Limit orders and the bid-ask spread. *Journal of Financial Economics* 53, 255-287.

Elmaghraby, Wedad, Altan Gulcu, and Pinar Keskinocak. 2005. Optimal markdown mechanisms in the presence of rational customers with multi-unit demands. Working paper, Georgia Institute of Technology.

Etzioni, O., C. Knoblock, R. Tuchinda, and A. Yates, 2003. To buy or not to buy: Mining airfare data to minimize ticket purchase price. *SIGKDD* ’03, August 24-27, 2003. Washington, DC, USA.

Foucault, T. 1999. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets* 2, 99-134.

Gallego, Guillermo, and Garrett van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizon. *Management Science* 40, 999-1020.

Harris, L. and J. Hasbrouck. 1996. Market vs. limit orders: the SuperDOT evidence on order submission strategy. *Journal of Financial and Quantitative Analysis* 31, 213-231.

Ho, Teck-Hua, Christopher S. Tang, and David R. Bell. 1998. Rational shopping behavior and the option value of variable pricing. *Management Science* 44, 145-160.

Johnson, Eric, Wendy Moe, Peter Fader, Steven Bellman, and Gerald Lohse. 2004. On the depth and dynamics of online search behavior. *Management Science* 50 (3), 299-308.

Kimes, Sheryl E. 1989. Yield management: a tool for capacity-constrained service firm. *Journal of Operations Management* 8, 348-363.

Kincaid, W. M., D. A. Darling. 1963. An inventory pricing problem. *Journal of Mathematical Analysis and Applications* 7 (2), 183–208.

Knapp, Louise. April 9, 2003. Algorithms key to cheap air fare. Wired News.

Liddle, Alan. 2003. Using web for discounting clicks with digital diners. *Nation's Restaurant News* 37 (20), 172.

Littlewood, K. 1972. Forecasting and control of passengers. *12th AGIFORS Symposium Proceedings*. Nathanya, Israel 95–128.

Liu, Qian and Garrett J. van Ryzin. 2005. Strategic capacity rationing to induce early purchases. Working paper, Columbia University.

McGill, Jeffrey, and Garrett van Ryzin. 1999. Revenue Management: Research Overview and Prospects. *Transportation Science* 33, 233-256.

Montgomery, Alan, Kartik Hosanagar, Ramayya Krishnan, and Karen Clay. 2004. Designing a better shopbot. *Management Science* 50 (2), 189-206.

Robinson, L. W. 1995. . Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research* 43 252-263.

Talluri, Kalyan, and Garrett J. van Ryzin . 2004. *The theory and practice of revenue management*. Kluwer Academic Publishers.

Zhao, Wen, and Yu-Sheng Zheng. 2000. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Science* 46, 375-388.